



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/760,345	01/12/2001	Kenji Yamagami	36992.00068	2223

30256 7590 12/21/2004

SQUIRE, SANDERS & DEMPSEY L.L.P
600 HANSEN WAY
PALO ALTO, CA 94304-1043

EXAMINER

KLINGER, SCOTT M

ART UNIT	PAPER NUMBER
----------	--------------

2153

DATE MAILED: 12/21/2004

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application N .

09/760,345

Applicant(s)

YAMAGAMI, KENJI

Examiner

Scott M. Klinger

Art Unit

2153

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 08 November 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-24 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-24 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application (PTO-152) |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

Claims 1-24 are pending.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 23 and 24 are rejected under 35 U.S.C. 102(e) as being anticipated by Carter et al. (U.S. Patent Number 6,553,401, hereinafter “Carter”). Carter discloses a system for implementing a high volume availability server cluster including both sharing volume of a mass storage on a local site and mirroring a shared volume on a remote site.

In referring to claim 23, Carter shows,

- Using a heartbeat check module to activate/de-activate a remote mirror and check for a heartbeat signal:

“A server 116A, 116B, ... 116Z may be unavailable to provide the database service for many reasons such as a hardware failure of the server, a software failure of the server, a power failure of the site at which the server is located, and/or a network failure preventing clients 102A, 102B, ... 102Z access to the server. In an exemplary embodiment, the cluster manager determines whether a server 116A, 116B, ... 116Z of the current subcluster 112A, 112B, ... 112Z is available based upon heartbeat signals transmitted amongst the servers 116A, 116B, ... 116Z of the server cluster 106.” (Carter, col. 6, lines 29-39)

Art Unit: 2153

A heartbeat check module that determines if network services are available and enables/disables said services accordingly is inherently implied in a system in which subclusters provide fail over.)

- The remote mirror being separate from the network:

Carter, Fig. 1 shows the mirrors **108A, 108B, 108Z** are separate from the network **114A, 114B, 14Z**

In referring to claim 24, Carter shows,

- Selectively activating a network between a primary group and a secondary group; selectively activating a remote mirror between the primary group and the secondary group:

Figure 2 shows a flow chart of the operation of the cluster system

- Checking for a failure occurrence in the primary group; if the network is activated, then sending a failure notification message from the primary group to the secondary group along the network; if the remote mirror is activated, then sending a failure notification message from the primary group to the secondary group along the remote mirror; based upon the failure notification message, displaying in the secondary group an indication of the failure occurrence (*U.S. Patent No. 6,553,401, col. 6, lines 29-39, quoted above*)

"It should be further appreciated that the distributed nature of the cluster server 106 and the mirroring of data resources between sites protects a service from a complete site failure. For example, if a database service is being provided by a first server 116A.sub.1 of a first subcluster 112A located in San Francisco, Calif. and a complete site failure occurs due to an earthquake in San Francisco, the cluster manager may reallocate the database service to a first server 116B.sub.1 of a second subcluster 112B located in Indianapolis, IN without an appreciable interruption of service." (Carter, col. 8, lines 4-13)

- The remote mirror being separate from the network:

Carter, Fig. 1 shows the mirrors **108A, 108B, 108Z** are separate from the network **114A, 114B, 14Z**

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1-8 and 10 are rejected under 35 U.S.C. 103(a) as being unpatentable over Ofek (U.S. Patent Number 6,044,444, hereinafter "Ofek") in view of Sicola et al. (U.S. Patent Number, 6,643,795, hereinafter "Sicola"). Ofek shows substantial features of the claimed invention, including:

- A first host group:
Ofek, Fig.4 shows a first host group **212**
- A first storage system associated with the first host group:
Ofek, Fig.4 shows a first storage system **214** associated with the first host group **212**
- The first host group coupled via a network to a second host group:
Ofek, Fig.4 shows the first host group **212** coupled via a network to a second host group **252**
- The first storage system coupled via a remote link separate from the network to a second storage system associated with the second host group:
Ofek, Fig.4 shows the first storage system **214** coupled via a remote link **241** separate from the network to a second storage system associated with the second host group
- The first host group configured to selectively send signals by use of the remote link by sending a memory command to the first storage system which is mirrored via the remote link using a remote copy mechanism to the second storage system:
"Accordingly, data may be transferred between the primary and secondary data storage system controllers synchronously, when a primary host computer requests writing of data to a primary data storage device, or asynchronously with the primary host computer requesting the writing of data to the primary data storage system, in which case the

Art Unit: 2153

remote data copying or mirroring is completely independent of and transparent to the host computer system.” (Ofek, col. 2, lines 50-57)

However, Ofek does not show sending a heartbeat signal. Nonetheless this feature is well known in the art and would have been an obvious modification to the system disclosed by Ofek.

In analogous art, Sicola discloses controller-based bi-directional remote copy system with storage site failover capability. Sicola shows sending a heartbeat signal: *“VA layer 520 is not aware of any PPRC manager 515 context (state change or transfer path). Host port target code 505 allows only host initiators to connect to the controller port which is a dedicated data replication port ... PPRC manager 515 is responsible for managing functions including initiating the connection and heartbeat with the remote controller and initiating the remote copy for incoming host writes (via host port initiator 510)”* (Sicola, col.9, lines 31-43)

Given these teachings, a person of ordinary skill in the art would have readily recognized the desirability and advantages of modifying the system of Ofek so as to use a copy mechanism for the heartbeat, such as taught by Sicola, in order to check if the disks are operational.

Claims 9 and 11-22 are rejected under 35 U.S.C. 103(a) as being unpatentable over Carter in view of Sicola.

In referring to claims 9 and 11, Carter shows substantial features of the claimed invention, including:

- A production host group; A standby host group coupled to the production host group by a network:

Figure 1 shows multiple host groups coupled by a network, elements 112A and 112B

- A remote mirror coupled between the production host group and the standby host group:
“Another step of the method includes mirroring the shared volume to a second mass storage device of a second subcluster that is located at a second site and that includes at least one server in order to obtain a first mirrored copy of the shared volume at the second site.” (Carter, col. 2, lines 22-27)

Art Unit: 2153

- The remote mirror including a production site heartbeat storage volume (heartbeat PVOL) and a standby site heartbeat storage volume (heartbeat SVOL) coupled by a remote link to the heartbeat PVOL:

Figure 1, Storage volumes 108A are connected to Storage volumes 108B are coupled through the network

- The production host group configured to selectively send a heartbeat signal to the standby host group by use of the remote link:

"In an exemplary embodiment, the cluster manager determines whether a server 116A, 116B, ... 116Z of the current subcluster 112A, 112B, ... 112Z is available based upon heartbeat signals transmitted amongst the servers 116A, 116B, ... 116Z of the server cluster 106." (Carter, col. 6, lines 34-39)

- The remote mirror being separate from the network:

Carter, Fig. 1 shows the mirrors **108A, 108B, 108Z** are separate from the network **114A, 114B, 14Z**

However, Carter is silent as to how the heartbeat signal is sent. Nonetheless this feature is well known in the art and would have been an obvious modification to the system disclosed by Carter as evidenced by Sicola. In analogous art, Sicola discloses controller-based bi-directional remote copy system with storage site failover capability. Sicola shows: *"VA layer 520 is not aware of any PPRC manager 515 context (state change or transfer path). Host port target code 505 allows only host initiators to connect to the controller port which is a dedicated data replication port ... PPRC manager 515 is responsible for managing functions including initiating the connection and heartbeat with the remote controller and initiating the remote copy for incoming host writes (via host port initiator 510)"* (Sicola, col.9, lines 31-43)

Given these teachings, a person of ordinary skill in the art would have readily recognized the desirability and advantages of modifying the system of Carter so as to use a copy mechanism for the heartbeat, such as taught by Sicola, in order to check if the disks are operational.

In referring to claim 12, Carter in view of Sicola discloses,

- A first heartbeat check module configured to generate the heartbeat signal:

Art Unit: 2153

A module configured to generate a heartbeat signal is inherently implied in a system that sends said heartbeat signal

In referring to claim 13, Carter in view of Sicola discloses,

- A second heartbeat check module configured to receive the heartbeat signal:

A module configured to receive a heartbeat signal is inherently implied in a system that receives said heartbeat signal

In referring to claim 14, Carter in view of Sicola discloses,

- The standby host group manages operations of the cluster computing system if an invalid heartbeat signal is received by the standby host group from the production host group:

The standby host group taking over operations if the production host group fails is inherently implied in a fail over system that utilizes a standby host group

“Moreover, the computer readable medium includes instructions, which when executed, cause a cluster manager to determine to reallocate the service to a first server of the second subcluster, allocate the first mirrored copy to the first server of the second subcluster, and allocate the service to the first server of the second subcluster in response to determining to reallocate the service to the first server of the second subcluster.”

(Carter, col. 3, lines 8-16)

In referring to claim 15, Carter in view of Sicola discloses,

- A serial number assigned to the heartbeat message; a time indicator indicating a time of the generation of the heartbeat message; an identifier identifying a sender of the heartbeat message:

Heartbeat messages (synonymous with “keepalive” packets) have an identifier (ip address), the time of generation, and a serial number by definition

Art Unit: 2153

In referring to claim 16, Carter in view of Sicola discloses,

- A second remote mirror coupled between the production host group and the standby host group:

“The cluster manager is operable to allocate the service and the at least one volume of the first mass storage device to a first server of the first subcluster, and mirror the at least one volume of the first mass storage device to the at least one volume of the second mass storage device.” (Carter, col. 2, lines 49-53)

Mirroring “at least” one volume inherently implies the mirroring of more than one volume is an embodiment of the invention.

- The second remote mirror including a second remote link for transmitting a heartbeat signal:

Figure 1, Storage volumes 108A are connected to Storage volumes 108B are coupled through the network

In referring to claim 17, Carter shows substantial features of the claimed invention, including:

- Generating a heartbeat signal from a production host group; selectively sending the heartbeat signal to the standby host group from the production host group by use of a remote link:

Carter, col. 6, lines 34-39 (See full quote above)

- Enabling the standby host group to manage operations of the cluster computing system if an invalid heartbeat signal is received by the standby host group from the production host group:

“Moreover, the cluster manager is operable to determine to reallocate the service to a first server of the second subcluster, allocate the at least one volume of the second mass storage device to a first server of the second subcluster, and allocate the service to the first server of the second subcluster in response to determining to reallocate the service to the first server of the second subcluster.” (Carter, col. 2, lines 53-60)

However, Carter is silent as to how the heartbeat signal is sent. Nonetheless this feature is

Art Unit: 2153

well known in the art and would have been an obvious modification to the system disclosed by Carter as evidenced by Sicola. In analogous art, Sicola discloses controller-based bi-directional remote copy system with storage site failover capability. Sicola shows: *"VA layer 520 is not aware of any PPRC manager 515 context (state change or transfer path). Host port target code 505 allows only host initiators to connect to the controller port which is a dedicated data replication port ... PPRC manager 515 is responsible for managing functions including initiating the connection and heartbeat with the remote controller and initiating the remote copy for incoming host writes (via host port initiator 510)"* (Sicola, col.9, lines 31-43)

Given these teachings, a person of ordinary skill in the art would have readily recognized the desirability and advantages of modifying the system of Carter so as to use a copy mechanism for the heartbeat, such as taught by Sicola, in order to check if the disks are operational.

In referring to claim 18, Carter in view of Sicola shows,

- Selectively sending a heartbeat signal to the production host group from the standby host group by use of a second remote link:

Carter, col. 6, lines 34-39 (See full quote above)

In referring to claim 19, Carter in view of Sicola shows,

- Registering a first storage volume to a device address entry, the first storage volume located in a production site, and, from the production site, changing a remote mirror that includes the first storage volume into an enabled mode; sending an activation message from the production site to a standby site; registering a second storage volume to the device address entry, the second storage volume located in the standby site; from the standby site, changing the remote mirror into an enabled mode to install a remote mirror formed by the first storage volume and second storage volume:

"Pursuant to another embodiment of the present invention, there is provided a server cluster for providing high availability of a service. The server cluster includes a first mass storage device located at a first site, a second mass storage device located at a second site, a first subcluster located at the first site, a second subcluster located at the

second site, and a cluster manager. The first mass storage device includes at least one volume associated with the service. Similarly, the second mass storage device includes at least one volume associated with the service. The first subcluster includes a plurality of servers operably coupled to the first mass storage device. Moreover, the second subcluster includes at least one server operably coupled to the second mass storage device. The cluster manager is operable to allocate the service and the at least one volume of the first mass storage device to a first server of the first subcluster, and mirror the at least one volume of the first mass storage device to the at least one volume of the second mass storage device. Moreover, the cluster manager is operable to determine to reallocate the service to a first server of the second subcluster, allocate the at least one volume of the second mass storage device to a first server of the second subcluster, and allocate the service to the first server of the second subcluster in response to determining to reallocate the service to the first server of the second subcluster” (Carter, col. 2, lines 36-60)

In referring to claim 20, although Carter in view of Sicola shows substantial features of the claimed invention including the method of checking for failure in a cluster computing system, Carter in view of Sicola does not explicitly show de-installing a remote mirror. Nonetheless this feature is well known in the art and would have been an obvious modification to the system disclosed by Carter in view of Sicola. Carter in view of Sicola discloses: *“It should be appreciated by those skilled in the art that mirroring is a continuous process. Accordingly, even though the flowchart of FIG. 2 illustrates mirroring as a distinct step of the exemplary operation, the mirror copy of the data resources stored on the second shared storage device 108B at site B is continuously being updated in order to reflect the current state of the data resources as stored on the first shared storage device 108A at site A.”* (Carter, col. 6, lines 4-12)

The mirrored copy of the data is continually updated, but Carter is silent as to how a failure of the mirror is handled. However, Carter in view of Sicola discloses that when the first subcluster fails or is unavailable, the resources are reallocated to the subcluster with the mirrored volume. Once the mirrored volume no longer exists (i.e. it becomes the primary volume), a new

Art Unit: 2153

mirror is created. Given these teachings, a person of ordinary skill in the art would have readily recognized the desirability and advantages of modifying the system of Carter in view of Sicola so as to de-install the failed mirror, and create a new mirror, as shown by the operation of the primary subcluster.

In referring to claim 21, Carter in view of Sicola shows,

- Sending heartbeat messages to the production site host if said production host is enabled; sending heartbeat messages to the standby site host if said standby host is enabled:

"A server 116A, 116B, ... 116Z may be unavailable to provide the database service for many reasons such as a hardware failure of the server, a software failure of the server, a power failure of the site at which the server is located, and/or a network failure preventing clients 102A, 102B, ... 102Z access to the server. In an exemplary embodiment, the cluster manager determines whether a server 116A, 116B, ... 116Z of the current subcluster 112A, 112B, ... 112Z is available based upon heartbeat signals transmitted amongst the servers 116A, 116B, ... 116Z of the server cluster 106." (Carter, col. 6, lines 29-39)

In referring to claim 22, Carter in view of Sicola shows,

- Checking for heartbeat messages from the production site host or the standby site host if the network is enabled, if an invalid heartbeat is received along the network and along the remote mirror, enabling the standby host to manage operations of the cluster computing system:

"Moreover, the cluster manager is operable to determine to reallocate the service to a first server of the second subcluster, allocate the at least one volume of the second mass storage device to a first server of the second subcluster, and allocate the service to the first server of the second subcluster in response to determining to reallocate the service to the first server of the second subcluster." (Carter, col. 2, lines 53-60)

Carter, col. 6, lines 29-39 (See full quote above)

Art Unit: 2153

Conclusion

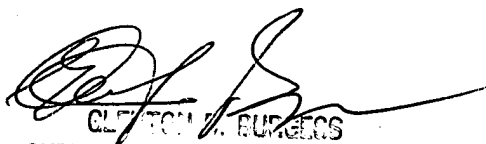
Any inquiry concerning this communication or earlier communications from the examiner should be directed to Scott M. Klinger whose telephone number is (703) 305-8285. The examiner can normally be reached on M-F 7:00am - 3:30pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Glenn Burgess can be reached on (703) 305-4792. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Scott M. Klinger
Examiner
Art Unit 2153

smk


GLENN W. BURGESS
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2153